

ATTORNEY DOCKET NO.
020431.0463

PATENT APPLICATION
09/156,334

Appendix C

XP-002090173

The Integrated Supply Chain Management System

P.D.	04.12.1993	11
P.	1-11	

Mark S. Fox, John F. Chionglo, Mihai Barbuceanu
Department of Industrial Engineering, University of Toronto

4 Taddle Creek Road, Toronto Ontario M5S 1A4

tel: 1-416-978-6823; fax: 1-416-978-3453; internet: msf@ie.utoronto.ca

Area: Manufacturing and Production Systems

Abstract: This paper describes the goals and architecture of the Integrated Supply Chain Management System (ISCM) being developed at the University of Toronto. ISCM provides an approach to the realtime performance of supply chain functions.

1.0 Introduction

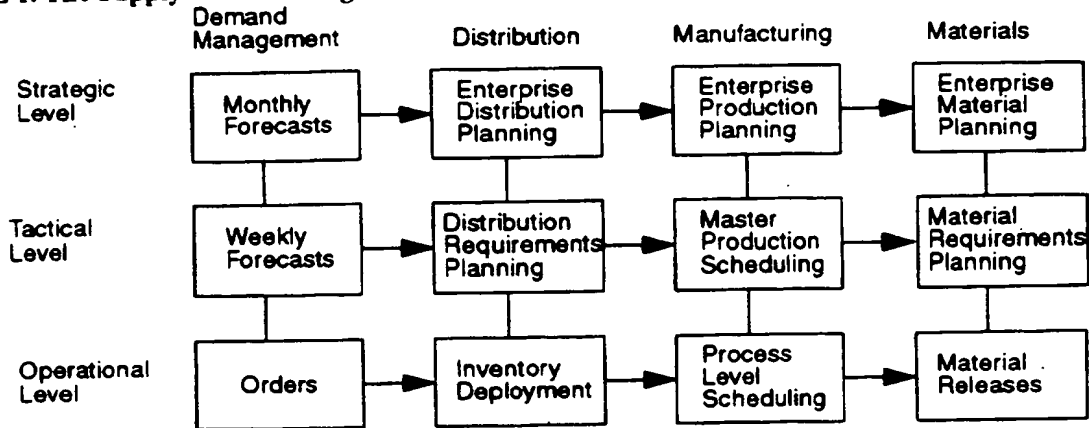
This paper describes the architecture of the Integrated Supply Chain Management System (ISCM) under development in the Enterprise Integration Laboratory at the University of Toronto.

In response to competitive pressures, managers are focusing on the reengineering of operations. Processes are being streamlined and automated, and work teams are reorganized and redeployed for higher productivity. Together with these changes, companies are looking for ways to better plan and control their operations. They are shifting away from a company with rigid and pre-planned activities to one that is able to react quickly and appropriately to changes.

The supply chain is a set of activities which span enterprise functions from the ordering and receipt of raw materials through the manufacturing of products through the distribution and delivery to the customer. In order to operate efficiently, these functions must operate in an integrated manner. Providing rapid and quality responses to supply chain events requires the coordination of multiple functions across the enterprise.

Supply chain management functions operate on three levels: strategic level, tactical level, and operational level.

FIGURE 1. The Supply Chain Management Functions



Each level is distinguished by the period of time over which decisions are made, and the granularity of decisions during that period. The strategic level addresses issues like: where to allocate production, and what is the best sourcing strategy. The tactical level addresses issues like: forecasting, scheduling, ordering of short lead time materials, and do we schedule overtime to meet production requirements. The operations level addresses issues like: inventory deployment, detailed scheduling, and what to do with an order when a machine breaks down.

Supply chain management also requires coordination with customers and suppliers. The dynamics of the market make this difficult. Customers often make changes or cancel orders. Suppliers may provide incorrect materials or deliver late. Systems that can quickly respond to market dynamics while minimizing lead times and inventory are required.

Like the market, the production floor is also dynamic. Unplanned events occur and cause deviations from scheduled activities. To achieve planned production, it is necessary for the production control system to dynamically respond to these events in ways that optimizes production goals. In some cases, events cause problems that are not "locally contained". The production control system must coordinate its actions with higher-level functions such as planning, sales, and marketing.

In the remainder of this paper, we describe the architecture of the ISCM system, its agents and their interactions.

2.0 System Design Issues

We view the supply chain as being managed by a set of intelligent (software) agents, each responsible for one or more activities in the supply chain, and each interacting with other agents in the planning and execution of their responsibilities. An agent is a software process that operates asynchronously, communicating with other agents as needed.

The first issue we face is deciding how supply chain activities should be distributed across the agents. Existing decompositions, as found in MRP systems, were limited by the sophistication, or lack thereof, of algorithms. This is exemplified by the distinction between MRP I and MRP II,

• which arises out of the move from infinite to finite Master Production Scheduling. We believe that the successful planning and execution of supply chain activities relies upon more sophisticated planning and scheduling algorithms than are available in current MRP systems. We view the planning/scheduling function as the "conductor" that "orchestrates" the behaviour of the other supply chain agents. Consequently, the nature of the reasoning performed by other agents will change. With more sophisticated planning/scheduling algorithms, the overall quality of supply chain management will increase.

The second issue is the nature of interactions among agent? Given the dynamics of the supply chain resulting from unplanned for (stochastic) events such as transportation problems, supply problems, etc., what is the nature of the interactions among agents that will result in the reduction of change-induced perturbations in a coordinated manner? If each agent has more than one way to respond to respond to an event, how do they cooperate in creating a mutually acceptable solution? In other words, how do agents influence or constrain each other's problem-solving behaviour?

The third issue is responsiveness. In a dynamic environment, the time available to respond may vary based on the event. It is a requirement that an agent's algorithm be able to respond within the time allotted. Algorithms that are able to generate a solutions no matter how much time is available are know as "anytime" algorithms. The quality of the solution of anytime algorithms is usually directly related to the time available.

The fourth issue is the availability of knowledge encapsulated within a module. In conventional MRP systems, a module is designed to perform a specific task. The modules may contain certain knowledge (used in the performance of each task) that could be used to answer related questions. It is our goal to "open up" a module's knowledge so that it can be used to answer questions beyond those originally intended.

In summary, the next generation supply chain management system will possess the following characteristics:

Distributed: The functions of supply chain management are divided among a set of separate, asynchronous software agents.

Dynamic: Each agent performs its functions asynchronously as required, as opposed to a batch or periodic mode.

Intelligent: Each agent is an "expert" in its function. Uses Artificial Intelligence and Operations Research problem solving methods.

Integrated: Each agent is aware of and can access the functional capabilities of other agents.

Responsive: Each agent is able to ask for information and/or a decision from another agent - each agent is both a client and a server.

Reactive: Each agent is able to respond to events as they occur modifying is behaviour as required, as opposed to responding in a pre-planned, rigid, batch approach.

Cooperative: Each agent can cooperate with other agents in finding a solution to a problem - that is, they do not act independently.

Interactive: Each agent may work with people to solve a problem.

Anytime: No matter how much time is available, an agent is able to respond to a request, but the quality of the response is proportional to the time given to respond.

Complete: The total functionality of the agents must span the range of functions required to manage the supply chain.

Reconfigurable: The supply chain management system itself must be adaptable and must support the "relevant subset" of software agents. For example, if the user only wants to schedule a plant, he/she should not be required to use or have a logistics component.

General: Each agent must be adaptable to as broad a set of domains as possible.

Adaptable: Agents need to quickly adapt to the changing needs of the human organization. For example, adding a resource or changing inventory policy should be quick and easy for the user to do.

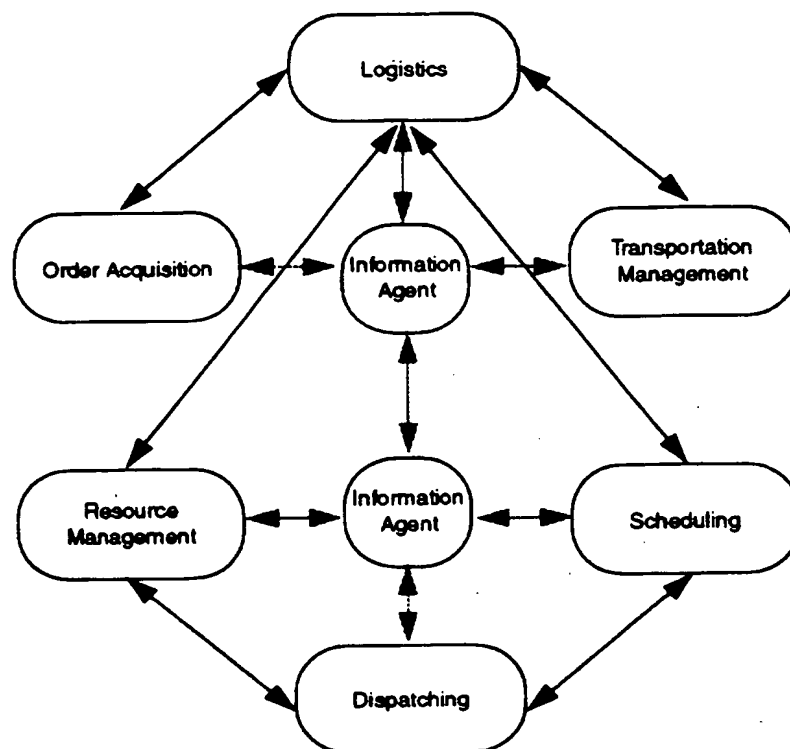
Backwards Compatible: Agents need to have a seamless upgrade path so that the release of new or changed features does not compromise existing integration or functionality.

3.0 Architectural Overview

The ISCM is composed of a set of cooperating agents, where each agent performs one or more supply chain management functions, and coordinates its decisions with other relevant agents. There are two types of agents: functional agents and information agents. Functional agents plan and/or control activities in the supply chain. Information agents support other agents by providing information and communication services.

The decomposition of supply chain functions and their allocation to agents represents one of the first tasks in the project. The problem is that existing decompositions of functions, as found in MRP systems today, arose out of organizational constraints, legacy systems and limitations on algorithms. For example, the distinction between Master Production Scheduling and Detailed Scheduling is primarily due to algorithm limitations. The merging of these two functions and the inclusion of some activities found in Inventory Management and Activity Planning is possible with the availability of more sophisticated planning and scheduling algorithms. We are currently working on six functional agents: Logistics, Transportation Management, Order Acquisition, Resource Management, Scheduling and Dispatching. They are described in more detail in the next section.

FIGURE 2. The ISCM agents



The dynamics of the environment make cooperative behaviour an important factor in integrating supply chain agents. In order to optimise supply chain decisions, an agent cannot make a locally optimal decision, but must determine the affect its decisions will have on other agents, and choose an alternative that optimises the entire supply chain

An agent is responsible for a set of functions or activities in the supply chain. Each agent stores information and knowledge locally and it may access information and knowledge throughout the network. We assume that the agents are in a heterogenous environment; hence, their interactions are made through message-based transactions.

Supply chain agents exist within an Enterprise Information Architecture (EIA). The EIA provides a distributed information environment where information may be stored anywhere in the network. The EIA manages the consistency of information. Subsets of information may be designated has being globally consistent and the EIA manages it. Other information, in which copies are stored locally by agents, may develop inconsistencies and are of no concern to the EIA.

The EIA provides each agent with automated information acquisition and distribution. When an agent requests information, the EIA will find it. When an agent creates information of interest to others, the EIA will distribute it to those agents that wish to know. The EIA provides the "right information in the right way" to the decision makers.

At the core of the EIA and the supply chain management system lies a generic reusable enterprise model. In order to support the integration of supply chain agents, it is necessary that shareable representation of knowledge be available that minimizes ambiguity and maximizes understanding and precision in communication. The enterprise model will also support "deductive query processing". Many of the terms in the generic model will be defined using Prolog axioms. These axioms will automate the answering of a significant number of questions raised by the system's users, thereby reducing software development costs.

4.0 Functional Agents

As said earlier, we believe that the successful planning and execution of supply chain activities relies upon more sophisticated planning and scheduling algorithms than are available in current MRP systems. We view the planning/scheduling function as the "conductor" that "orchestrates" the behavior of the other supply chain agents. Consequently, the nature of the reasoning performed by other agents will change. With more sophisticated planning/scheduling algorithms, the overall quality of supply chain management will increase. The rest of this section describes briefly each of the functional agents under development.

Order Acquisition agent. This agent is responsible for acquiring orders from customers, negotiating with customers about prices, due dates, etc., and handling customer requests for modifying or canceling respective orders. This agent is one of the agents participating in negotiation that may be necessary to successfully create supply chain plans. These will be exceptional situations where other agents find an over-constrained situation requiring modification of constraints.

This agent captures the order information from directly from customers and communicates these orders to the logistics agent. When a customer order is changed, it is communicated to the logistics agent. When plans violate constraints imposed by the customer (such as due date violation), the order acquisition agent participates in negotiating with the customer and the logistics agent for a feasible plan.

Logistics agent. This agent is responsible for coordinating multiple-plants, multiple-supplier, and multiple-distribution center domain of the enterprise to achieve the best possible results in terms of goals of the supply chain, which include ontime delivery, cost minimization, etc. It manages the movement of products or materials across the supply chain from the supplier of raw materials to the customer of finished goods.

The inputs to the logistics agent are customer orders, deviations in factory schedules which affects customer orders, transportation plans and resource availabilities. The outputs of the agent are production requirements for each factory, supplier, etc., and transportation requirements.

Transportation agent: This agent is responsible for the assignment and scheduling of transportation resources in order to satisfy inter-plant movement requests specified by the Logistics Agent. It will be able to consider a variety of transportation assets and transportation routes in the construction of its schedules.

Scheduling agent. This agent is responsible for scheduling and rescheduling activities in the factory, exploring hypothetical "what-if" scenarios for potential new orders, and generating schedules that are sent to the dispatching agent for execution.

The inputs to the scheduling agent are production requests from the logistics agent, resource problems from the resource agent, and the deviations of the current schedule from the dispatching agent. Its output is a detailed schedule.

The scheduling agent assigns resources and start times to activities that are feasible while at the same time optimizing certain criteria such as minimizing WIP or tardiness. It can generate a schedule from scratch or repair an existing schedule that has violated some constraints.

In anticipation of domain uncertainties such as machine breakdowns, material inavailability, etc., the agent may reduce the precision of a schedule by increasing the degrees of freedom in the schedule for the dispatcher to work with. For example, it may "temporally pad" a schedule by increasing an activity's duration, or "resource pad" an operation by either providing a choice of more than one resource or increasing the capacity required so that more is available.

The scheduling agent also acts as a coordinator when infeasible situations arise. It has the capability to explore tradeoffs among the various constraints and goals that exist in the plant.

Resource agent. The resource agent merges the functions of inventory management and purchasing. It dynamically manages the availability of resources so that the schedule can be executed. It estimates resource demand and determines resource order quantities. It is responsible for selecting suppliers that minimizes costs and maximizes delivery. It generates EDI purchase requests and monitors their fulfillment.

The inputs to the resource agent are the schedule from the scheduler, the availability or unavailability of resources from suppliers, the arrival of resources from the factory floor, and the consumption of resources from the dispatcher. The outputs of the resource agent include the arrival of resources, the availability of resources, and the orders sent to suppliers.

The resource agent generates purchase orders and monitors the delivery of resources. When resources do not arrive as expected, it assists the scheduler in exploring alternatives to the schedule by generating alternative resource plans.

Dispatching agent. This agent performs the order release and realtime floor control functions as directed by the scheduling agent. It operates autonomously as long as the factory performs within the constraints specified by the scheduling agent. When deviations from schedule occur, the dispatching agent communicates them to the scheduling agent for repair.

The inputs to the dispatching agent are the schedule from the scheduling agent, the status of the factory floor, and the availability of resources. The outputs are the deviations from the current schedule and the starting of activities.

Given degrees of freedom in the schedule, the dispatcher makes decisions as to what to do next. In deciding what to do next, the dispatcher must balance the cost of performing the activities, the

amount of time in performing the activities, and the uncertainty of the factory floor. For example, a) given that the scheduler specified a time interval for the start time of a task, the dispatcher has the option of either starting the task as soon as possible (JIC) or starting the task as late as possible (JIT), b) given that the scheduler did not specify a particular machine for performing the task, the dispatcher may use the most "cost effective" machine (minimize costs) or use the "fastest" machine (minimize processing time).

5.0 Enterprise Information Architecture

An Enterprise Information Architecture (EIA) provides communication and information services supporting:

- Persistent storage of information to be shared among the multiple functional agents in the corporate network.
- Deductive capabilities allowing new information to be inferred from existing information.
- Automatic distribution of information to the agents that need it.
- Automatic retrieval, processing and integration of information that is relevant to agents.
- Checking and maintaining various forms of consistency of the information.
- Performing information access control functions such as determining who is allowed to see and change the available information.

The EIA is composed of both functional agents and information agents (IA). An IA services a number of agents (functional and other IA-s) by providing them with a layer of shared information storage and services for managing it. Agents periodically volunteer some of their information to the IA (and keep it up to date) or just answer the queries sent to them by the IA. The IA uses its own information together with the supplied information to determine which information needs of other agents can be satisfied. It processes the information in order to determine the most relevant content and the most appropriate form for the needs of these agents. In the process, it may uncover various forms of inconsistency among the supplied information and take action to remove them. IA-s will also communicate with each other in order to accomplish their functions.

IA-s are not meant to replace the direct communication channels established among agents during their usual interactions. Rather, they support these interactions by providing shared access to information and the basic information management services listed above. IA-s will be particularly useful in cases when:

1. A consistent form of shared information needs to be maintained.
2. Information from many sources needs to be aggregated, perhaps in a continuous fashion, to produce reports or answer queries.
3. Information has to be distributed among many agents.

4. Changes in the state of the modeled enterprise need to be propagated over the models and activities of various agents.

5. Inconsistencies arising during agent interaction need to be quickly detected and resolved.

6.0 Agent Interaction

Given the dynamics of the supply chain resulting from unplanned for (stochastic) events such as transportation problems, supply problems, etc., what is the nature of the interactions among agents that will result in the reduction of change-induced perturbations in a coordinated manner? If each agent has more than one way to respond to an event, how do they cooperate in creating a mutually acceptable solution? In other words, how do agents influence or constrain each other's problem-solving behavior?

In order for two or more agents to cooperate, there must exist a "cultural assumption". The existence of a cultural assumption implies what an agent can expect in terms of another agent's behavior in a problem solving situation. A possible cultural assumption is that agents are "constraint-based problem solvers." That is, given a set of goals and constraints, they search for a solution that optimizes the goals and satisfies the constraints. Another cultural assumption could be that agents have the ability to generate more than one solution. Thereby the enabling the consideration of alternatives and trade-offs by a set of cooperating agents. A third cultural assumption is that agents have the ability and authority to suboptimize local goals and possibly relax a sub set of constraints if the global solution is further optimized.

Our approach to coordination is to view agent problem-solving as a constraint satisfaction/optimization process. An agent solves a problem by first understanding what constraints and goals exist, then intelligently searching for a solution that satisfies the constraints and optimizes the goals as best as it can. When an agent's problem-solving relies upon or affects the problem-solving of another agent, it must interact with it. We believe that an agent can modify another's problem-solving behavior through the communication of constraints. Research has demonstrated the power of this approach [Fox 83] [Fox 86] [Fox & Sycara 90] [Sycara et al. 92]. Coordination occurs when agents develop plans that satisfy their own internal constraints but also the constraints of other agents. Negotiation occurs when constraints, that cannot be satisfied, are modified by the subset of agents directly concerned.

7.0 Enterprise Model

At the core of the EIA and the supply chain management system lies a generic reusable enterprise model. In order to support the integration of supply chain agents, it is necessary that shareable representation of knowledge be available that minimizes ambiguity and maximizes understanding and precision in communication. The goal of the TOVE Enterprise Modelling project is to create a data model that has the following characteristics: 1) provides a shared terminology for the enterprise that each agent can jointly understand and use, 2) defines the meaning of each term (aka semantics) in a precise and as unambiguous manner as possible, 3) implements the semantics in a set of axioms that will enable TOVE to automatically deduce the answer to many "common

sense" questions about the enterprise, and 4) defines a symbology for depicting a term or the concept constructed thereof in a graphical context.

We approach the first goal by defining a generic level representation which the application representations are defined in terms of. Generic concepts include representations of Time [Allen 84], Causality [Rieger 77] [Bobrow 85], Activity [Sathi 85], and Constraints [Fox 83] [Davis 87]. The generic level is, in turn, defined in terms of a conceptual level based on the 'terminological logic' of KLONE [Brachman 85].

We approach the second and third goals by defining a set of axioms (aka rules) that define common-sense meanings for the terminology. By common sense, we mean that the more obvious definitions/deductions about the entities and attributes in our ontology. (We view definitions as being mostly circuitous, as opposed to be reducible to a small set of grounded terms.) What is an obvious deduction should be determined by a subset of questions used to determine the competence of a representation. Since there does not exist a standard for determining the competence of a model, we will define, in english, a set of questions and the axioms used to answer them.

To date we have developed ontologies for activity, state, time, causality, resources, constraints, quality, cost and organization structure.

TOVE is not only a research project but a testbed. TOVE has been used to implement a *virtual company* whose purpose is to provide a testbed for research into enterprise integration. TOVE is implemented in C++ using the ROCK@+[TM] knowledge representation tool from Carnegie Group. TOVE operates "virtually" by means of knowledge-based simulation [Fox 89].

8.0 Conclusions

The goals of the Integrated Supply Chain Management Project are:

- Identifying an appropriate decomposition of supply chain functions and encapsulate into agents.
- Developing protocols and strategies for the communication of information, coordination of decisions, and management of change.
- Develop/use state-of-the-art algorithms for agent decision-making.
- Developing an incremental, "anytime" model of problem solving for each functional agent so that it can provide rapid responses to unplanned for events.
- Extending each function oriented agent so that it is able to answer more questions within its functional domain.

9.0 References

- [Allen 83] Allen, J.F. Maintaining Knowledge about Temporal Intervals. *Communications of the ACM*. 26(11):832-843. 1983.
- [Allen 84] Allen, J.F. Towards a General Theory of Action and Time. *Artificial Intelligence*. 23(2):123-154. 1984.
- [Bobrow 85] Bobrow, D.G. *Qualitative Reasoning About Physical Systems*. MIT Press, 1985.
- [Bobrow 77] Bobrow, D., and Winograd, T. KRL: Knowledge Representation Language. *Cognitive Science*. 1(1). 1977.
- [Brachman 77] Brachman, R.J. *A Structural Paradigm for Representing Knowledge*. PhD thesis, Harvard University, 1977.
- [Brachman 79] Brachman, R.J. On the Epistemological Status of Semantic Networks. *Associative Networks: Representation and Use of Knowledge by Computers*. In Findler, N.V., Academic Press, 1979, pages 3-50.
- [Brachman 85] Brachman, R.J., and Schmolze, J.G. An Overview of the KL-ONE Knowledge Representation Systems. *Cognitive Science*. 9(2), 1985.
- [Davis 87] Davis, E. Constraint Propagation with Interval Labels. *Artificial Intelligence*. 3281-331. 1987.
- [Esprit 90] ESPRIT-AMICE. CIM-OSA - A Vendor Independent CIM Architecture. *Proceedings of CINCOM 90*, pages 177-196. National Institute for Standards and Technology, 1990.
- [Fahlman 77] Fahlman, S.E. *A System for Representing and Using Real-World Knowledge*. PhD thesis, Massachusetts Institute of Technology, 1977.
- [Fox 79] Fox, M.S. On Inheritance in Knowledge Representation. *Proceedings of the International Joint Conference on Artificial Intelligence*. 95 First St., Los Altos, CA 94022, 1979.
- [Fox 81] Fox, M.S. An Organizational View of Distributed Systems. *IEEE Transactions on Systems, Man, and Cybernetics*. SMC-11(1):70-80. 1981.
- [Fox 83] Fox, M.S. *Constraint-Directed Search: A Case Study of Job-Shop Scheduling*. PhD thesis, Carnegie Mellon University, 1983. CMU-RI-TR-85-7, Intelligent Systems Laboratory, The Robotics Institute, Pittsburgh, PA.
- [Fox 89] Fox, M.S., Reddy, Y.V., Husain, N., McRoberts, M. Knowledge Based Simulation: An Artificial Intelligence Approach to System Modeling and Automating the Simulation Life Cycle. *Artificial Intelligence, Simulation and Modeling*. In Widman, L.E., John Wiley & Sons. 1989.